Service Guarantees for Joint Scheduling and Flow Control [1]

R. L. Cruz (*cruz@ece.ucsd.edu*) C. M. Okino (*cokino@ece.ucsd.edu*)
Department of Electrical & Computer Engineering
University of California, San Diego
La Jolla, CA 92093-0407

Abstract

We consider a general mesh network with multiple traffic streams subject to window flow control on a per hop, per stream basis. Scheduling at each server is governed by "service curve" requirements. We establish lower bounds on the window sizes such that each stream receives pre-specified service guarantees.

Keywords: Quality of Service Guarantees, Scheduling, Statistical Multiplexing, Window Flow Control

1 Introduction

The allocation of network resources to multiple streams in an integrated services network in conjunction with a flow control scheme is a recent topic of interest in the literature. Open loop rate-based flow control (e.g. "leaky bucket") has been proposed at network access points [19] in high speed integrated services networks in order to eliminate or reduce buffer overflow, providing a mechanism for the network to support quality of service guarantees. Under certain conditions, imposing open loop rate-based flow control within the network does not increase worst case end-to-end delay [20] [8] [10], and in fact can increase the capacity of networks that operate with certain scheduling algorithms [11]. However, the use of open loop rate-based flow control may increase *average* end-to-end delay, since the flow of traffic is sometimes inhibited even when resources are underutilized. The implication is that, in certain cases, open loop rate-based flow control may prevent efficient statistical multiplexing. However, it remains attractive in high speed networks because of its insensitivity to large bandwidth-delay products.

Window-based flow control [3] can directly eliminate the possibility of buffer overflow in communication networks by relying on feedback. It has the potential for providing for efficient statistical multiplexing. For networks with large bandwidth delay products, a large window size and possibly large buffers may be required for efficient bandwidth utilization. To partially cope with large bandwidth-delay products, hop-by-hop window flow control has been proposed (also known as "credit-based" flow control in ATM networks , e.g. see [15]). It is commonly believed that window-based flow control cannot provide any performance guarantees. Perhaps for this reason, window-based flow control is often thought to be appropriate only for delay tolerant applications (e.g. Available Bit Rate (ABR) traffic). However, we believe that, in conjunction with rate-based flow control at access points and appropriate scheduling algorithms within the network, window-based flow control *can* provide performance guarantees and may offer an attractive alternative to open loop rate-based flow control within the network.

In an earlier paper [9], we studied the flow of traffic on a *single* stream subject to window flow control, where cross traffic was modelled through constraints on servers. In this paper, we explicitly model *multiple* streams in a general mesh network, where each stream is subject to window flow control and packet scheduling is governed by "service curve" requirements, as defined in Section 2.

We take a deterministic approach along the lines of [6], since we believe this is a productive way to develop initial insight into complex queueing phenomena. Indeed, feedback that arises in window flow control results in complex interactions between queues, which are often difficult to analyze rigorously with stochastic models.

The first mathematically rigorous bound on the throughput of window flow control protocols was apparently first developed by E. L. Hahne [12], in a data network setting. Hahne derived lower bounds on window sizes which insured that throughputs were close to max-min fair [3], assuming round-robin scheduling at each network node. Our work is similar to Hahne's in that we also derive lower bounds on window sizes so that pre-specified service requirements are met.

In the following section, we introduce the concept of "service curves". In Section 3, we describe our model in detail. We present our main result in Section 4, and conclude in Section 5 with some brief remarks.

2 Background

In this paper, we draw heavily upon the general concept of service curves [8], which has roots in the work of Parekh and Gallager [16]. The service curve concept will be used to describe the model we define in Section 3.

2.1 Service Curves

We consider a discrete time model, where time is divided into slots, numbered $0, 1, 2, \cdots$. Consider a net-

[1] This work was supported by the National Science Foundation under grants NCR-91-58618 and NCR 94-15684, by the Army Research Office under grant FRI DAAH04-95-1-0248, and by the Air Force Office of Scientific Research under grant F49620-95-1-0538

work element with entering and exiting traffic described by the rate functions r^{in} and r^{out}. Specifically, we define $r^{in}[t]$ to be the number of packets arriving in slot t and $R^{in}(t)$ to be the number of packets arriving in the interval $[1,t]$. We similarly define $r^{out}[t]$ to be the number of packets departing from the network element in slot t and $R^{out}(t)$ to be the number of packets departing from the network element in the interval $[1,t]$. The number of packets stored in the network element at the end of slot $t \geq 0$ is $B[t] = R^{in}(t) - R^{out}(t)$, where we assume $B[0] = 0$. Suppose S is a given non-negative function. To simplify the notation, assume without loss of generality that $S(x) = 0$ for all $x \leq 0$.

Definition 1. (Service Curve Guarantee). *A system is said to guarantee the service curve S if for all $t \geq 0$, there exists $s \leq t$ such that $R^{out}(t) - R^{in}(s) \geq S(t-s)$.*

Definition 1 has been previously reported in [17]. It has also been independently made by Agrawal and Rajan [1] and Le Boudec [14]. Earlier related definitions of service guarantees were made independently by Cruz [7], Hung and Kesidis [13], and Stiliadis and Varma [18].

Given two functions F and G defined on the non-negative integers, define the convolution of F and G, written $F * G$, as

$$F * G(x) \quad = \min_{x_1+x_2=x, x_1,x_2 \geq 0} \{F(x_1) + G(x_2)\} \ .$$

The convolution operator is analogous to conventional convolution prevalent in linear system theory, but is in terms of the "min-plus algebra". In the min-plus algebra, the minimum operation is associated with the usual addition operation, and the addition ("plus") operation is associated with the usual multiplication operation. The minimum and "plus" operation form an algebra in the sense that they are both commutative and associative operators, and "plus" is distributive over "min." For more information and applications, the reader is referred to [2]. It is easy to verify that the convolution operation is commutative and associative, and that it distributes over the minimum operation. It is straightforward to verify that the following is in fact equivalent to Definition 1.

Alternative Definition 1. (Service Curve Guarantee). *A system is said to guarantee the service curve S if for all $t \geq 0$ there holds $R^{out}(t) \geq S(t) * R^{in}(t)$.*

Define the "impulse function" $\hat{\delta}(x) = 0$ if $x \leq 0$, and $\hat{\delta}(x) = +\infty$ if $x > 0$. Note that for any function F, $F * \hat{\delta}(x) = F(x)$. It is interesting to note then that a service curve is analogous in some sense to the concept of an impulse response of a linear time invariant system. The connection to linear system theory has also been independently reported by Chang [4].

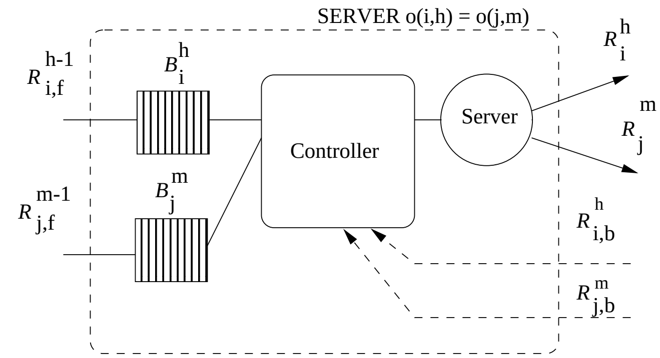

Figure 1: Server $o(i,h)$ with streams i and j.

The virtual delay at slot t, $D[t]$, is defined as

$$D[t] = \min\{\Delta : \Delta \geq 0 \text{ and } R^{out}(t+\Delta) \geq R^{in}(t)\} \ .$$

Theorem A. [8] *Assume that R_{in} satisfies the following "burstiness constraint": $R^{in}(t) - R^{in}(s) \leq b(t-s)$ for all $s \leq t$. Suppose a system guarantees a service curve of S. Then*
(a) (Buffer Requirements) There holds for all t

$$B[t] \leq \max_{\alpha \geq 0}\{[b(\alpha) - S(\alpha)]^+\} \ .$$

(b) (Bound on Delay) There holds for all t

$$D[t] \leq \max_{\alpha \geq 0}\{\min\{\Delta : \Delta \geq 0 \ \textit{and} \ b(\alpha) \leq S(\alpha+\Delta)\}\} \ .$$

(c) (Output Burstiness) For all $s \leq t$ there holds $R^{out}(t) - R^{out}(s) \leq b^{out}(t-s)$, where

$$b^{out}(x) = \max_{\Delta \geq 0}\{b(x+\Delta) - S(\Delta)\} \ .$$

Theorem B. [8] (Convolution Theorem) *Consider traffic flowing through a system consisting of n subsystems in tandem, where the j^{th} subsystem guarantees the service curve S^j. Then the system as a whole guarantees the service curve $S^{net} = S^1 * S^2 * \cdots * S^n$.*

3 Network Model

A packet switched network is represented by a directed graph $G = (V, E)$ where V is the set of vertices and E is the set of edges. A vertex in this graph represents a packet switch, and an edge in this graph represents a unidirectional communication link. Within each packet switch, there is a queueing server associated with each outgoing link that schedules packet transmissions on that link.

For simplicity, we assume that each server has the same capacity c. Specifically, the maximum number of packets that a server can serve is assumed to be c per slot, for each server in G. We index the servers by the integers.

We assume that a finite set of traffic streams is flowing through the network of servers, and we index these traffic streams by the integers. Stream i flows through a fixed sequence of n_i servers. Let $o(i,h)$ be the index of the server visited by stream i at hop h. Thus, stream i traverses a path in the graph, such that $\{o(i,1), o(i,2), \ldots, o(i,n_i)\}$ is the ordered sequence of servers visited by traffic stream i. Define, $I_{i,h}$ to be the set of traffic streams which share server $o(i,h)$, i.e. $I_{i,h} = \{(j,m) : o(j,m) = o(i,h)\}$. Similarly, let I_e be the set of traffic streams which pass through the server for link e in G.

Let $r_i^h[t]$ and $R_i^h(t)$ be the number of packets departing server $o(i,h)$ from stream i during slot t and during the interval $[1,t]$, respectively (i.e. $R_i^h(t) = \sum_{n=1}^{t} r_i^h[n] = r_i^h[1,t]$). The packets departing server $o(i,h)$ pass through network element $N_{i,f}^h$ before reaching server $o(i,h+1)$. For example, network element $N_{i,f}^h$ may model *forward* propagation delay. We define $r_{i,f}^h[t]$ and $R_{i,f}^h(t)$ as the number of packets entering server $o(i,h+1)$ (and hence departing network element $N_{i,f}^h$) from stream i during slot t and during the interval $[1,t]$, respectively. For notational convenience, let $r_{i,f}^0[t]$ and $R_{i,f}^0(t)$ be the number of packets entering the network (i.e., arriving to server $o(i,1)$) from stream i during slot t, and during the interval $[1,t]$, respectively. Each stream i entering server $o(i,h+1)$ first enters its own buffer (see Figure 1). Remembering that packets arrive at server $o(i,h+1)$ at rate $r_{i,f}^h$ and assuming that the network is empty at slot 0, the amount of packets from the i^{th} stream held at server $o(i,h+1)$ at the end of slot t is

$$B_i^{h+1}[t] = R_{i,f}^h(t) - R_i^{h+1}(t) \ . \tag{1}$$

For all i and h we assume that a "target service curve" $\hat{S}_i^h$ is given. We would like stream i to be guaranteed the service curve $\hat{S}_i^h$ at server $o(i,h)$, i.e. we would like the following inequality to hold for all t:

$$R_i^h(t) \geq R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \tag{2}$$

Toward this end, packets from each stream are assigned deadlines in the following manner. The packets from stream i are indexed in order of arrival, starting from index 1. For example, if exactly one packet arrives from stream i to server $o(i,h)$ during slot t, it has the index $R_{i,f}^{h-1}(t)$. Packet k arriving to server $o(i,h)$ from stream i in slot t is assigned the deadline $D_{i,k}^h$ where

$$D_{i,k}^h = \min\{\Delta : \Delta \geq 0 \text{ and } R_{i,f}^{h-1} * \hat{S}_i^h(\Delta) \geq k\} \ . \tag{3}$$

Intuitively, these deadlines are assigned so that, if met, the departure process $R_i^h(t)$ will not cross below $R_{i,f}^{h-1} * \hat{S}_i^h(t)$. The deadline assignments in (3) are consistent with the so-called scheduling "SCED" policy defined in [17]. The following lemma demonstrates that it is possible to compute these deadlines at server $o(i,h)$ in real-time. In other words, if packet k from stream i arrives to the server $o(i,h)$ at time t, then the deadline for that packet can be calculated without knowledge of $R_{i,f}^{h-1}(s)$ for $s > t$.

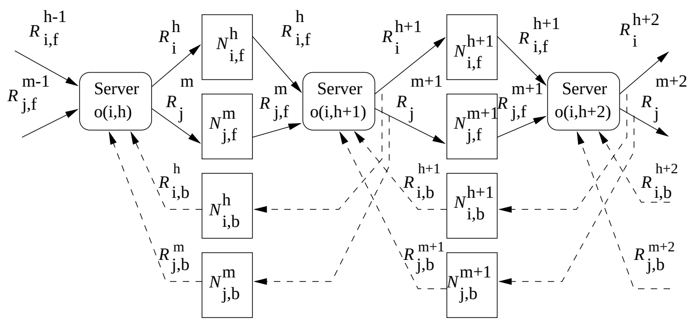

Figure 2: Example of three servers along stream i.

Lemma 1 *If the k^{th} packet of stream i arrives to server $o(i,h)$ at time t, then $D_{i,k}^h = D_{i,k}^h(t)$, where*

$$D_{i,k}^h(t) \quad = \quad \min\{u : u \geq t \text{ and } g(u) \geq k\} \ , \tag{4}$$

and $g(u) = \min_{s:0 \leq s \leq t-1}\{R_{i,f}^{h-1}(s) + S_i^h(u-s)\}$.

The next lemma will be useful later.

Lemma 2 *The number of packets from stream i that are assigned deadlines which lie in $[1,t]$, for the server $o(i,h)$, is equal to $Z_i^h(t)$, where*

$$Z_i^h(t) = R_{i,f}^{h-1} * S_i^h(t) \tag{5}$$

for $h = 1, 2, \cdots, n_i$.

3.1 Scheduling Algorithm

Nominally, the server $o(i,h)$ serves packets in an earliest deadline first manner. More specifically, in each slot t the server serves packets with the earliest deadlines among those "available" and "eligible." A packet is available in slot t if it is queued in the (buffer at the) server from the previous slot, or it arrives to server $o(i,h)$ during slot t (i.e. we consider a "cut-through" mode of operation). A packets eligibility is determined by a window flow control protocol.

If $h = n_i$, then an available packet at server $o(i,h)$ is always eligible. In other words, a stream is not subject to window flow control at its last hop. Otherwise, server $o(i,h)$ and server $o(i,h+1)$ exchange "tokens" to determine if packets at server $o(i,h)$ from stream i are eligible. There are K_i^h tokens initially available at server $o(i,h)$ for stream i, where K_i^h is a non-negative integer, called the "window size." Each packet from stream i must acquire a token in order to become eligible. A packet from stream i departing from server $o(i,h)$ carries the acquired token with it. After the packet receives service from server $o(i,h+1)$, the token is sent back to server $o(i,h)$ through a network element $N_{i,b}^h$. More precisely, the number of tokens sent back to server $o(i,h)$ (and hence entering network element $N_{i,b}^h$) from server $o(i,h+1)$ during slot t is $r_i^{h+1}[t]$. Let $r_{i,b}^h[t]$ (resp. $R_{i,b}^h(t)$) be the number of tokens from stream i returning to server $o(i,h)$ from server $o(i,h+1)$ (and hence departing network element $N_{i,b}^h$) during slot t (resp. during the interval $[1,t]$). For example, the network element $N_{i,b}^h$ could model *backward* propagation delay.

Figure 2 is an example of three servers in tandem along the path of stream i. We assume in this specific example that stream j follows the same path as stream i for at least these three hops, i.e. $o(i,h) = o(j,m)$, $o(i,h+1) = o(j,m+1)$, and $o(i,h+2) = o(j,m+2)$. Note that in general, each stream need not take the same path as any other stream for any number of hops.

The number of outstanding tokens (or the number of "unacknowledged packets") for stream i at server $o(i,h)$ at the end of slot t is defined to be $T_i^h[t]$, where

$$T_i^h[t] = R_i^h(t) - R_{i,b}^h(t) \ . \tag{6}$$

If $T_i^h[t-1] < K_i^h$, then during slot t there is at least one token available at server $o(i,h)$ for packets from stream i, and any available packet may acquire it. If $T_i^h[t-1] = K_i^h$ then no tokens are available for stream i at the beginning of slot t; However, if $r_{i,b}^h[t] > 0$, at least one token returns from server $o(i,h+1)$ during slot t, and an available packet from stream i at server $o(i,h)$ may aquire it (and hence potentially immediately depart server $o(i,h)$ during slot t.) If $T_i^h[t-1] = K_i^h$ and $r_{i,b}^h[t] = 0$, then no tokens are available for stream i at server $o(i,h)$, and stream i is said to be *blocked* at server $o(i,h)$ during slot t. In this case, note that $T_i^h[t] = K_i^h$.

In summary, in each slot t, server $o(i,h)$ serves as many available packets as it can, up to a maximum of c packets; The packets are chosen among all streams which are not blocked, such that packets with earlier deadlines receive priority for service.

For notational convenience, if $h = n_i$, we define $R_i^{h+1} \equiv R_{i,b}^h \equiv R_i^h$, $S_i^{h+1} = \hat{\delta}_0$, and $K_i^h = \infty$. Thus, (2) is trivially true for $h = n_i + 1$.

We characterize network elements $N_{j,f}^m$ and $N_{j,b}^m$ by "service curves" $S_{j,f}^m$ and $S_{j,b}^m$, respectively for all $e \in E$ and $(j,m) \in I_e$. Specifically, we assume that

$$R_{i,f}^h(t) \geq R_i^h * S_{i,f}^h(t) \tag{7}$$

and

$$R_{i,b}^h(t) \geq R_i^{h+1} * S_{i,b}^h(t) \tag{8}$$

for all t. In addition, we assume that

$$S_{j,b}^m(1) = 0 \ , \tag{9}$$

for all $e \in E$ and $(j,m) \in I_e$.

Finally, we assume the resource allocation condition

$$\sum_{(j,m)\in I_e} \hat{S}_j^m(t) \quad \leq \quad ct \ , \tag{10}$$

for all $e \in E$ and $t = 0, 1, 2 \ldots$.

4 Main Result

We are interested in analyzing traffic for an arbitrary stream i and an arbitrary hop h. If we can find a service curve guarantee for stream i at server $o(i,h)$, then an end-to-end service curve guarantee for any stream in the network can be determined by the convolution theorem, and bounds on end-to-end delay are implied by Theorem A.

Note that if blocking never occurs in the network then effectively there is no window flow control and each stream (i,h) is guaranteed the service curve $\hat{S}_i^h$ at server $o(i,h)$ [8]. Further, note that no matter how large the window size K_i^h may be, blocking of stream (i,h) is always possible since we assume no upper or lower bounds to the amount of traffic each source generates. Despite possible blocking by window flow control, our main result, Theorem 1, establishes that *finite* window sizes are sufficient to guarantee that stream (i,h) receives service curve $\hat{S}_i^h$ for all (i,h).

Before stating the main result, we introduce some convenient notation.

For all $e \in E$ let

$$\hat{K}_j^m \quad = \quad \max_{x\geq 0}\{\hat{S}_j^m(x) - \hat{S}_j^m * S_{j,loop}^m(x)\} + \Delta_j^m \quad \text{for all } (j,m) \in I_e \ ,$$

where

$$\Delta_j^m = \max_{x\geq 0}\{\hat{S}_j^m(x) - \hat{S}_j^m * \hat{S}_j^m(x)\}$$

and

$$S_{j,loop}^m(x) \quad = \quad S_{j,b}^m * \hat{S}_j^{m+1} * S_{j.f}^m(x) \quad \text{for all } (j,m) \in I_e.$$

Theorem 1. Service Curve Guarantee *If (7) - (10) hold and the window sizes satisfy $K_j^m \geq \hat{K}_j^m$ for all $(j,m) \in I_e$ then for any $(i,h) \in I_e$, server $o(i,h)$ guarantees stream (i,h) the service curve $\hat{S}_i^h$, i.e. for all t there holds*

$$R_i^h(t) \quad \geq \quad R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \tag{11}$$

Furthermore,

$$R_i^h(t) \quad \geq \quad R_{i,f}^{h-1} * \hat{S}_i^h(t) + \Delta_i^h, \text{ if } T_i^h[t] = K_i^h. \tag{12}$$

To prove Theorem 1, we make use of the following lemmas.

Lemma 3 *Suppose (7) - (9) hold and $(j,m) \in I_e$ for any $e \in E$. Suppose also that the total amount of unacknowledged traffic on stream j satisfies $T_j^m[p] = K_j^m$, and $R_j^{m+1}(u) \geq R_{j,f}^m * \hat{S}_j^{m+1}(u)$ for all $u < p$. Then there exist a slot $q < p$ such that the total amount of packets served by server $o(i,h)$ on stream j over the interval $[q+1,p]$ satisfies*

$$R_j^m(p) - R_j^m(q) \quad \geq \quad S_{j,loop}^m(p-q) + K_j^m \ . \tag{13}$$

Furthermore, we have,

$$p - q \quad \geq \quad \frac{K_j^m}{c} \ . \tag{14}$$

Lemma 4 *Suppose $(i,h) \in I_e$ and $e \in E$, and fix t. Let*

$$\tau \quad = \quad \max\{s : s \leq t \textit{ and } \sum_{(j,m)\in I_e} r_j^m[s] < c\}$$

If all packets served in the interval $[\tau, t]$ at server $o(i,h)$ have deadlines greater than t, then define $p = \tau$. Otherwise let p be the last slot in $[\tau, t]$ when a packet with deadline greater than t is served. If $I^*_{i,h}$ is the set of streams where server $o(i,h)$ has served at least one packet from that stream in the interval $[p+1, t]$ and $(j,m) \in I^*_{i,h}$ such that $B_j^m[p] > 0$, then

$$T_j^m[p] \quad = \quad K_j^m \ . \tag{15}$$

Lemma 5 *Suppose (7) holds and $(j,m) \in I_e$ for any $e \in E$. Suppose further that the total number of packets served by server $o(i,h)$ on stream $(j,m) \in I_{i,h}$ over the interval $[1,q]$ satisfies*

$$R_j^m(q) \quad \leq \quad R_{j,f}^{m-1} * \hat{S}_j^m(q) \ .$$

If p is such that $p \leq q$ and either

(a) $B_j^m[p] = 0$

or

(b) $T_j^m[p] = K_j^m$ *and* $R_j^m(p) \geq R_{j,f}^{m-1} * \hat{S}_j^m(p) + \Delta_j^m$,

then

$$R_j^m(q) - R_j^m(p) \quad \leq \quad \hat{S}_j^m(q-p) \ . \tag{16}$$

Lemma 6 *Fix t and let $(i,h) \in I_e$, $e \in E$, and assume that (7) - (10) hold. Suppose also that for all $s \leq t-1$, $(j,m) \in I_{i,h}$, we have $R_j^m(s) \geq R_{j,f}^{m-1} * \hat{S}_j^m(s)$ and that $R_j^m(s) \geq R_{j,f}^{m-1} * \hat{S}_j^m(s) + \Delta_j^m$ if $T_j^m[s] = K_j^m$.*

(i) If $T_i^h[t] = K_i^h$, then there exist $t_1 < t$ such that

$$R_i^h(t) - R_i^h(t_1) \geq S_{i,loop}^h(t - t_1) + K_i^h \ .$$

(ii) If $R_i^h(t) < R_{i,f}^{h-1} * \hat{S}_i^h(t)$ and $T_i^h[t] < K_i^h$, then there exist $s^* < t$ such that

$$R_i^h(t) - R_i^h(s^*) \geq \hat{S}_i^h(t - s^*) \ ,$$

and if $B_i^h[s^*] > 0$, there exist $t_1 < s^*$ such that

$$R_i^h(s^*) - R_i^h(t_1) \geq S_{i,loop}^h(s^* - t_1) + K_i^h \ .$$

Proof of Theorem 1: We will use time induction to prove the theorem. Fix any slot $t \geq 0$ and let $(i,h) \in I_e$.

Let $\mathcal{H}(t)$ be the induction hypothesis that for all $(j,m) \in I_e$, $e \in E$, and for all time $s \leq t-1$ there holds

$$R_j^m(s) \quad \geq \quad R_{j,f}^{m-1} * \hat{S}_j^m(s) \ , \tag{17}$$

and

$$R_j^m(s) \quad \geq \quad R_{j,f}^{m-1} * \hat{S}_j^m(s) + \Delta_j^m \quad , \text{ if } T_j^m[s] = K_j^m. \tag{18}$$

For $t = 1$, we have

$$R_i^h(0) = 0 \geq 0 = R_{i,f}^{h-1} * \hat{S}_i^h(0) \ ,$$

and furthermore, we have $T_i^h[0] = 0 < K_i^h$. Thus, $\mathcal{H}(1)$ holds. Assuming that $\mathcal{H}(t)$ holds, we now establish $\mathcal{H}(t+1)$.

Fix $(i,h) \in I_e$ for any $e \in E$. Suppose that $T_i^h[t] = K_i^h$. By the induction hypothesis $\mathcal{H}(t)$, for $t_1 < t$, we have $R_i^h(t_1) \geq R_{i,f}^{h-1} * \hat{S}_i^h(t_1)$. Using Lemma 6 (i), there exists $t_1 < t$ such that

$$\begin{aligned}
R_i^h(t) &= R_i^h(t_1) + R_i^h(t) - R_i^h(t_1) \\
&\geq R_i^h(t_1) + S_{i,loop}^h(t-t_1) + K_i^h \\
&\geq R_i^h(t_1) + S_{i,loop}^h(t-t_1) + \hat{K}_i^h \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t_1) + S_{i,loop}^h(t-t_1) + \hat{K}_i^h \\
&= R_{i,f}^{h-1}(v) + \hat{S}_i^h(t_1 - v) + S_{i,loop}^h(t-t_1) + \hat{K}_i^h \\
&\geq R_{i,f}^{h-1}(v) + \hat{S}_i^h * S_{i,loop}^h(t-v) + \hat{K}_i^h \\
&\geq R_{i,f}^{h-1}(v) + \hat{S}_i^h * S_{i,loop}^h(t-v) \\
&\quad + \hat{S}_i^h(t-v) - \hat{S}_i^h * S_{i,loop}^h(t-v) + \Delta_i^h \\
&\geq R_{i,f}^{h-1}(v) + \hat{S}_i^h(t-v) + \Delta_i^h \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t) + \Delta_i^h \qquad (19) \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \qquad (20)
\end{aligned}$$

Thus, (19) proves (18) for the hypothesis $\mathcal{H}(t+1)$.

Now suppose that $T_i^h[t] < K_i^h$. Furthermore, suppose that

$$R_i^h(t) \quad < \quad R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \tag{21}$$

By Lemma 6 (ii), there may exist $s^* < t$ such that $B_i^h[s^*] = 0$ and $R_i^h(t) - R_i^h(s^*) \geq \hat{S}_i^h(t - s^*)$, which gives us

$$\begin{aligned}
R_i^h(t) &= R_i^h(s^*) + R_i^h(t) - R_i^h(s^*) \\
&= R_{i,f}^{h-1}(s^*) + R_i^h(t) - R_i^h(s^*) \\
&\geq R_{i,f}^{h-1}(s^*) + \hat{S}_i^h(t - s^*) \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \qquad (22)
\end{aligned}$$

Otherwise, there exists $t_1 < s^* < t$ such that $B_i^h[s^*] > 0$ and $R_i^h(t_1) \geq R_{i,f}^{h-1} * \hat{S}_i^h(t_1)$, using $\mathcal{H}(t)$. Applying Lemma 6 (ii), this gives us

$$\begin{aligned}
R_i^h(t) &= R_i^h(t_1) + R_i^h(s^*) - R_i^h(t_1) \\
&\quad + R_i^h(t) - R_i^h(s^*) \\
&\geq R_i^h(t_1) + S_{i,loop}^h(s^* - t_1) + K_i^h \\
&\quad + \hat{S}_i^h(t - s^*) \\
&\geq R_i^h(t_1) + S_{i,loop}^h(s^* - t_1) + \hat{K}_i^h \\
&\quad + \hat{S}_i^h(t - s^*) \\
&\geq R_i^h(t_1) + \hat{S}_i^h * S_{i,loop}^h(t - t_1) + \hat{K}_i^h \\
&\geq R_i^h(t_1) + \hat{S}_i^h * S_{i,loop}^h(t - t_1) + \hat{S}_i^h(t - t_1) \\
&\quad - \hat{S}_i^h * S_{i,loop}^h(t - t_1) + \Delta_i^h
\end{aligned}$$

$$
\begin{aligned}
&= R_i^h(t_1) + \hat{S}_i^h(t-t_1) + \Delta_i^h \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t_1) + \hat{S}_i^h(t-t_1) + \Delta_i^h \\
&= R_{i,f}^{h-1}(v) + \hat{S}_i^h(t_1-v) + \hat{S}_i^h(t-t_1) + \Delta_i^h \\
&\geq R_{i,f}^{h-1}(v) + \hat{S}_i^h * \hat{S}_i^h(t-v) + \hat{S}_i^h(t-v) \\
&\quad - \hat{S}_i^h * \hat{S}_i^h(t-v) \\
&= R_{i,f}^{h-1}(v) + \hat{S}_i^h(t-v) \\
&\geq R_{i,f}^{h-1} * \hat{S}_i^h(t) \ . \qquad (23)
\end{aligned}
$$

Thus, (22) and (23) contradict (21), and so

$$R_i^h(t) \quad \geq \quad R_{i,f}^{h-1} * \hat{S}_i^h(t), \text{ if } T_i^h[t] < K_i^h. \qquad (24)$$

Finally, (20) and (24) prove (17) for the hypothesis $\mathcal{H}(t+1)$. $\diamond$

5 Closing Remarks

Given the results of [9][1][4][5], one might expect that a window size of $\hat{K}_{i,old}^h$ is sufficient to guarantee that stream (i,h) is guaranteed the service curve $\hat{S}_i^h$, where

$$\hat{K}_{i,old}^h = \max_x \{\hat{S}_i^h(x) - \hat{S}_i^h * \hat{S}_i^h * S_{i,loop}^h(x)\} \ . \qquad (25)$$

Comparing this to Theorem 1, it can be shown that $\hat{K}_{i,old}^h \leq \hat{K}_i^h$. The potentially larger window size $\hat{K}_i^h$ was apparently necessary here to insure that (16) holds.

The identification of guaranteed service curves for window sizes smaller than that specified in Theorem 1 is a subject for future investigation.

It would also be of interest to investigate window flow control protocols exercised on an aggregate basis, rather than a per-connection basis as in this paper.

References

[1] R. Agrawal and R. Rajan, "Performance Bounds for Guaranteed and Adaptive Services," preprint, Nov. 1996.

[2] F. L. Baccelli, G. Cohen, G. Oslder, and J. Quadrat, *Synchronization and linearity : an algebra for discrete event systems,* Wiley, 1992.

[3] D. P. Bertsekas and R. G. Gallager, *Data Networks,* 2nd ed., Prentice Hall, 1992.

[4] C.-S. Chang, "On deterministic traffic regulation and service guarantee: a systematic approach by filtering," preprint, to appear in IEEE INFOCOM'97.

[5] C.-S. Chang, "Matrix extensions of the filtering theory for deterministic traffic regulation and service guarantees," preprint, March 1997.

[6] R. L. Cruz, "A calculus for network delay, part I: network elements in isolation," *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 114-131, Jan. 1991.

[7] R. L. Cruz, "Service burstiness and dynamic burstiness measures: a framework," *Journal of High Speed Networks*, vol. 1, no. 2, pp. 105-127, 1992.

[8] R. L. Cruz, "Quality of service guarantees in virtual circuit switched networks," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 1048-1056, Aug. 1995.

[9] R. L. Cruz and C. M. Okino, "Service Guarantees for Window Flow Control," Proc. 34^{th} Allerton Conf. on Comm., Cont. & Comp., Monticello, IL, Oct. 1996.

[10] L. Georgiadis, R. Guerin, V. Peris, and R. Rajan, "Efficient support of delay and rate guarantees in an internet," Proc. ACM SIGCOMM'96, pp. 106-116, Aug. 1996.

[11] L. Georgiadis, R. Guerin, V. Peris, and K. N. Sivarajan, "Efficient network QoS provisioning based on per node traffic shaping," *IEEE/ACM Transactions on Networking,* vol. 4, no. 4, pp. 482-501, Aug. 1996.

[12] E. L. Hahne, "Round robin scheduling for fair flow control in data communications networks," *IEEE JSAC,* vol. 9, no. 7, pp. 1024-1039, Sept. 1991.

[13] A. Hung and G. Kesidis, "Bandwidth scheduling for wide-area ATM networks using virtual finishing times," *IEEE/ACM Transactions on Networking,* vol. 4., no. 1, pp. 49-54, Feb. 1996.

[14] J.-Y. Le Boudec, "Network calculus made easy," preprint.

[15] H. T. Kung, T. Blackwell, and A. Chapman, "Credit update protocol for flow controlled ATM networks: statistical multiplexing and adaptive credit allocation," Proceedings of SIGCOMM'94, pp. 101-115.

[16] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the multiple node case," *IEEE/ACM Trans. Networking,"* vol. 2, no. 2, pp. 137-150, April 1994.

[17] H. Sariowan, "A service-curve approach to performance guarantees in integrated-service networks," Ph.D. thesis, Dept. of Electrical & Computer Engineering, UCSD, June 1996.

[18] D. Stiliadis and A. Varma, "Latency-rate servers: a general model for analysis of traffic scheduling algorithms," Proc. IEEE INFOCOM'96, pp. 111-119, Mar. 1996.

[19] Turner, J., "New Directions in Communications (or Which Way to the Information Age?)," IEEE Communications Magazine, Vol 24., No 10, October 1986.

[20] H. Zhang and D. Ferrari, "Rate controlled service disciplines," *Journal of High Speed Networks,* vol.3 no. 4, pp. 389-412, 1994.